

WHAT IS CLAIMED IS:

1. A system for responding to file system requests having file IDs comprising V, a volume identifier specifying the file system being accessed, and R, an integer, specifying the file within the file system being accessed comprising:

D disk elements in which files are stored, where D is greater than or equal to 2 and is an integer;

a switching fabric connected to the D disk elements to route requests to a corresponding disk element; and

N network elements connected to the switching fabric, each network element having a mapping function that for every value of V, specifies one or more elements from the set D that store the data specified by volume V, where N is greater than or equal to 2 and is an integer and N + D is greater than or equal to 4, which receives the requests and causes the switching fabric to route the requests by their file ID according to the mapping function.

2. A system as described in Claim 1 wherein each network element includes a translator which obtains file IDs from path names included in individual file system requests.

3. A system as described in Claim 2 wherein each disk element and each network element has a file system location database which maintains a mapping from all file system identifiers V to disk element identifiers so each network element can translate

each file system request ID into a corresponding disk element location.

4. A system as described in Claim 3 wherein each disk element and each network element has a controller, and each disk element controller communicates with the network element controllers to identify which files are stored at the respective disk element.

5. A system as described in Claim 4 wherein each network element can respond to any request for any disk element.

6. A system as described in Claim 5 wherein each network element has a network port through which requests are received by the respective network element wherein all the network elements and disk elements together appear as a single system that can respond to any request at any network port of any network element.

7. A system as described in Claim 6 wherein the disk elements form a cluster, with one of the disk elements being a cluster coordinator which communicates with each disk element in the cluster to collect from and distribute to the network elements which file systems are stored in each disk element of the cluster at predetermined times.

8. A system as described in Claim 7 wherein the cluster coordinator determines if each disk element is operating properly and redistributes requests for any disk element that is not operating properly; and allocates virtual network interfaces to

network elements and assigns responsibility for the virtual network interfaces to network elements for a failed network element.

9. A system as described in Claim 8 wherein network elements and disk elements can be added dynamically.

10. A system as described in Claim 9 wherein each network element advertises the virtual interfaces it supports to all disk elements.

11. A system as described in Claim 10 wherein each disk element has all files with the same file system ID for one or more values of V.

12. A system as described in Claim 11 wherein each request has an active disk element and a passive disk element associated with each request, wherein if the active disk element fails, the passive disk element is used to respond to the request.

13. A system as described in Claim 11 wherein the requests include NFS requests.

14. A system as described in Claim 13 wherein the requests include CIFS requests.

15. A system as described in Claim 14 wherein the translator obtains the file IDs from path names contained within CIFS requests.

16. A method for responding to file system requests comprising the steps of:

receiving file system requests having file IDs comprising V, a volume identifier specifying the file system being accessed, and R, an integer, specifying the file within the file system being accessed at network elements, each having a mapping function that for every value of V, specifies one or more elements from the set D that store the data specified by volume V; and

routing the requests to a switching fabric connected to the network elements based on the file system request's ID according to the mapping function to disk elements connected to the switching fabric.

17. A method as described in Claim 16 wherein the receiving step includes the step of obtaining the ID from path names included in the requests with a translator of the network element.

18. A method as described in Claim 17 wherein the routing step includes the step of maintaining all disk element locations at each file system location database of each disk element and each network element so each network element can translate each file system request ID into a corresponding disk element location.

19. A method as described in Claim 18 wherein the receiving step includes the step of receiving requests at a network port of the network element which can respond to any request, and

all the network elements and disk elements together appear as a single system.

20. A method as described in Claim 19 wherein the routing step includes the step of collecting from and distributing to the disk elements and the network elements, which form a cluster, which file systems are stored in each disk element by a cluster coordinator, which is one of the disk elements of the cluster, at predetermined times.

21. A method as described in Claim 20 wherein the routing step includes the step of redistributing requests from any disk elements which are not operating properly to disk elements which are operating properly by the network elements which receive the requests.

22. A method as described in Claim 21 wherein after the routing step, there is the step of adding dynamically network elements and disk elements to the cluster so the cluster appears as one server and any host connected to any network port can access any file located on any disk element.

23. A method as described in Claim 22 wherein before the receiving step, there is the step of advertising by each network element each virtual interface it supports.

24. A method as described in Claim 23 wherein the obtaining step includes the step of obtaining ID requests by the translator of the network element from path names contained in a CIFS request.

25. A system for responding to file system requests having file IDs comprising V, a volume identifier specifying the file system being accessed, and R, an integer, specifying the file within the file system being accessed comprising:

D disk elements in which files are stored, where D is greater than or equal to 2 and is an integer;

a switching fabric connected to the D disk elements to route requests to a corresponding disk element; and

N network elements connected to the switching fabric, each network element having a mapping function that for every value of V, specifies one or more elements from the set D that store the data specified by volume V, where N is greater than or equal to 2 and is an integer and $N + D$ is greater than or equal to 4, wherein network elements and disk elements can be added dynamically.

26. A system for responding to file system requests having file IDs comprising V, a volume identifier specifying the file system being accessed, and R, an integer, specifying the file within the file system being accessed comprising:

D disk elements in which files are stored, where D is greater than or equal to 2 and is an integer;

a switching fabric connected to the D disk elements to route requests to a corresponding disk element; and

N network elements connected to the switching fabric, each network element having a mapping function that for every value of V, specifies one or more elements from the set D that store the data specified by volume V, where N is greater than or equal to 2 and is an integer and N + D is greater than or equal to 4, wherein each network element has a network port through which requests are received by the respective network element wherein all the network elements and disk elements together appear as a single system that can respond to any request at any network port of any network element.